REMARKS/ARGUMENTS

Favorable reconsideration of this application as presently amended and in light of the following discussion is respectfully requested.

Claims 1-43, 45-67 and 69-90 are pending, Claims 1, 2, 4-5, 14, 17, 22, 25, 43, 54, 55, 57, 62, 67, 79 and 81 having been amended and Claims 44 and 68 having been canceled by way of the present amendment. Applicant reserves the right to prosecute canceled Claims 44 and 68 in one or more continuation applications.

In the outstanding Office Action the drawings were objected to; Claims 1-90 were rejected under 35 U.S.C. §101; Claims 1-6, 8, 11, 22-27, 29, 32, 43-48, 50, 56, 67-72, 74 and 80 are rejected as being anticipated by Altschuler et al. (U.S. Patent 6,556,983, hereinafter Altschuler); and Claims 7, 9-10, 12-21, 28, 30-31, 33-42, 49, 51-55, 57-66, 73, 75-79, and 81-90 were rejected as being unpatentable over Altschuler in view of Oles ("An Application of Lattice Theory to Knowledge Representation").

The undersigned appreciatively acknowledges the courtesy extended by Examiner Pham in holding an interview with the undersigned on August 17, 2004. During the interview, the present invention was discussed and distinguished from the asserted prior art of Altschuler. While the undersigned suggested the prosecution should focus on independent Claims 43 and 67, which are directed to generating a finite sheaf (see specification beginning at page 15, line 14 to page 16, line 2 for example), it is believed that it would be more beneficial to provide a more complete explanation of the invention in the context of all of the claims, including Claims 1-22, which is done in the remarks below.

In reply to the objection to the drawings, a separate letter requesting entry of substitute drawings is filed herewith, changing Figure 4C by adding the element 90.

As discussed in the interview, the independent claims have been amended to address the 35 U.S.C. §101 rejection. It is respectfully submitted that the claims as amended, comply

with 35 U.S.C. §101. However, if the Examiner disagrees the Examiner is invited to telephone the undersigned so that mutually agreeable claim language may be identified.

Prior to an analysis of the asserted prior art, a more complete appreciation of the context in which the present invention relates is believed to be in order.

The present inventor was the first to identify two limitations of the conventional relational data model approach to computer-based data management: the "order limitation" and the "schema limitation". In the following discussion, these limitations and how the present invention removes them are described.

As is well known, the central data storage object in the relational model is the mathematical relation, more concretely thought of as a table. One applies the relational approach to a problem domain by modeling the domain as a family of entity types; each entity type consisting of a set of entities (instances) and each entity having specific values for each of the attributes specified for the type. The relational approach proceeds by creating a data storage object (a table) for each entity type; creating a column in the table for each attribute of the type; creating a row in the table for each entity of an entity type; and entering data in the cells of each row to represent the specific values of the attributes of the corresponding entity. The mathematical operations defined by the relational model can then be used to store, query, and update the data representing the problem domain.

The first limitation is the order limitation. It is an explicit feature of the relational data model that the rows of the table are not ordered. There is no explicit specification of a way to represent or store an order for the rows, and the operations of the relational model neither rely on nor maintain any notion of order to the rows. Yet, as recognized by the present inventor, there are important classes of applications that either require, or would benefit substantially from, being able to specify and maintain an order to the rows. Furthermore, the kind of order these applications need is typically not a simple sequential order, but a more complex order

associated with containment or inclusion in some sense. When using the relational approach, such applications are forced to either repeatedly re-compute ordering relationships from data that can be stored in the rows of the table; or to somehow force the ordering relationships into a form that can be stored in a separate table; or to store the ordering information outside the context of the relational model. All three of these options are complex, inefficient, and place an undesirable burden on the application.

The present invention addresses the order limitation by introducing a different mathematical class as the central data storage object. Instead of the relation, the invention uses the partially ordered set, which can be concretely thought of as a table plus a graph with a node for each row in the table. As can be seen in the several figures of the patent application, one can picture the graph as being attached to the rows of the table and, for want of a better term, the following description refers to this combination as a "row-graph-table". In a manner analogous to the conventional relational approach, one may apply the present invention to a problem domain by creating such a row-graph-table for each entity type, creating a row for each entity, and entering attribute data in each row. But in addition, one must also specify the ordering information, that is, the links in the graph part of the row-graph-table. The mathematical operations defined by the present invention can then be used to store, query and update both the row data and the graph (ordering) data simultaneously. Not only do these operations relieve the application from the burden of maintaining the ordering data, the extended capabilities they provide, relative to the conventional relational operators, can be used to simplify the work of the application in many ways.

Claims 1 and 42, as amended correspond to the discussion thus far. The above discussion applies to Claims 43, 45-67 and 69-90, as well, but the following discussion regarding the second limitation, the schema limitation, also relates to these higher-numbered claims. The relational model treats the schema of the data, that is, the specification of the

columns of the tables, as "metadata". The prefix "meta" literally means "along with" or "beyond" and metadata is data that is conceptually part of the database, but lies beyond the scope of the central data object and its associated operators. Conventionally, the tables and their columns are specified externally by a "database administrator" and the operators of the relational data model operate in the context of this fixed schema. The rows in the tables can be created, deleted and updated by the relational operators, but the columns are fixed. This limitation specifically precludes the case in which the columns of a given table depend on the rows in another table. Yet there is at least one large and important class of application data that requires precisely this capability, namely, spatial dependence data.

It is often the case that an application involves the dependence of some property on location within some spatial object. A few simple examples include:

- population versus location within a state
- temperature versus location within a building
- air flow velocity versus location on the surface of an airplane wing

In such cases, it is common practice to represent the spatial object as some collection of simpler parts and store each part as a row in a table. For instance:

- a state may be represented by a table with a row for each zip code
- a building may be represented by a table with a row for each room
- an airplane wing may be represented by a table with a row for each 1 inch by 1 inch square on its surface

If the spatial object is represented by a table with rows for its parts, then the property data is most naturally represented by a table with a column for each row in the spatial object table. Each row in such a table then represents the dependence of the property on location at some given time or under some given conditions. For instance:

- For population versus location, each column represents the population of a specific zip code. Each row represents population versus location at a specific time, perhaps one row for each census year.

- For temperature versus location, each column represents the temperature in a specific room. Each row represents temperature versus location at a specific time of day, perhaps one row for each hour.
- For velocity versus location, each column represents the air speed in some particular square inch on the wing surface. Each row represents velocity versus location under some specific conditions, perhaps at different climb rates.

The schema limitation makes it very difficult to represent this kind of data using the relational approach. Once again, the application must either somehow force the spatial dependence data into a form that meets this limitation or store the data outside the context of the relational model. Once again, these options are complex, inefficient, and place an undesirable burden on the application.

The present invention as defined by Claims 43-90 addresses this limitation with the same strategy as it addresses the order limitation: it uses a different mathematical class as the data storage type. In fact, it uses a mathematical class that incorporates the partially ordered set described above and thus removes both the ordering limitation and the schema limitation at the same time. This mathematical class is the sheaf (which is first introduced in Claim 43), which one can think of concretely as a table with two graphs, a row graph and a column graph, the column graph being the row graph of another table. By extension of the previous nomenclature (and for the purpose of this discussion only) this structure is called a "column-row-graph-table". The use of this object is still analogous to the conventional relational approach. One applies the sheaf model to a problem domain by creating a column-row-graph-table for each entity type, creating a row for each entity, and entering attribute data in each row. But in addition, when creating the table one must specify which other table to use for the column graph; and when entering row data, one must also specify the ordering information, the links in the row graph part of the column-row-graph-table. The mathematical operations defined by the sheaf model can then be used to store, query and update the row

data, the ordering data, and the column data simultaneously. Once again, these operations not only relieve the application from the burden of maintaining the ordering and column data, the extended capabilities they provide, relative to the conventional relational operators, can be used to simplify the work of the application in many ways.

Now that the background/context of the invention has been discussed, the discussion turns to a comparison of the claims and the asserted prior art. Altschuler is asserted in the outstanding Office Action as anticipating Claim 1. Applicant respectfully traverses the rejection.

Altschuler is directed to a system for using lattice structures in which data are modeled as graphs having "entity" nodes which depict properties or attributes of the data and "relationship" links which denote relationship among the entities (column 13, lines 33-36). Various types of data may be represented using annotated entity relationship diagrams (or "a-ERDs"). (column 13, lines 39-44). Figure 47 is a table data structure 4700 which represents usage log graph 800 of Figure 8. The data used in Altschuler are represented as graphs in an extended sense, in that relationship links can attach to other relationship links (as shown in Figure 8).

Comparing amended Claim 1 to Altschuler, amended Claim 1 is directed to constructing a distinct table to represent each entity type such that there is a one-to-one correspondence between tables and entity types. For each table, a column in the table is constructed for each attribute of the entity type such that there is a one-to-one correspondence between columns and respective attributes of the entity type. Each row in the table is constructed such that there is a row in the table for each entity of the entity type such that there is a one-to-one correspondence between rows and respective entities. A row graph is then constructed such there is a specified ordering relationship between the rows of the table.

Claim 1 further requires defining a finite distributive lattice from the row graph to have a distinct member for each row and each combination of the rows of the first table.

In contrast, Altschuler is not directed to a system that uses different tables for each entity type, and within each table, columns arranged by attribute and separate rows for each entity of the single entity type. Rather, Altschuler is directed to a system that operates on annotated entity relationship diagrams (a-ERD) that comprise an entire table of data (e.g., Figure 47 as an example of an a-ERD). Altschuler does not disclose defining a finite distributive lattice from a row graph to have a member for each row and for each distinct combination of rows of the table. Thus, while amended Claim 1 is directed to specific features within a table, Altschuler is directed to something entirely different, namely providing relationships between different annotated entity relationship diagrams. Furthermore, Claim 1 requires that the row graph represent an externally specified ordering relationship between the rows of the first table. In contrast, a feature of Altschuler is that the relationship between different a-ERDs is determined automatically and not externally specified.

Consequently, it is respectfully submitted that Claim 1, as amended, is patentably distinct over Altschuler. Although Claims 2-6, 8, 11, 22-27, 29, 32, 43-48, 50, 56, 67-72, 74 and 80, are of differing scope and/or statutory subject class, it is respectfully submitted that each of these claims also patentably defines over Altschuler for at least the same reasons discussed above with regard to Claim 1.

Turning to the rejections of pending claims over the combination of Altschuler in view of Oles, each of these rejections is based on the assertion of Altschuler as applied to the independent claims discussed above. Oles is asserted for its alleged description of defining a set of algebraic operators to operate on graphs, and lattice data models. However, Oles does not disclose the features of Claim 1 as discussed above that are absent in Altschuler , namely

the step of defining a finite distributive lattice from the row graph to have a member for each row and each distinct combination of the rows of the first table, as an example. As a consequence, no matter how Oles is combined with Altschuler, the combination would neither teach nor suggest all of the features of independent Claim 1, and therefore the combination does not teach or suggest the features of dependent Claims 7-10, 12-21, 28, 30-31, 33-42, 49, 51-55, 57-66, 73, 75-79, and 81-90.

Amended Claim 43 is directed to a computer implemented method of representing data on a computer as a finite sheaf. The steps, among other things, include constructing a first table to represent the entity type with a column in the table for a respective attribute of the entity type. This entity type is a single entity type as was discussed during the interview. The method also includes entering attribute data into rows of the first table. This first table includes only one row for each primary entity of the entity type such that there is a one-two-one correspondence between rows in respective primary entities. The method then includes a step of specifying a row graph which represents an externally specified ordering relationship of the rows of the first table. The method further includes defining a finite distributive lattice (FDL) from the row graph to have a member for each row and each distinct combination of the rows of the first table. As was discussed during the interview, the method of Claim 43 also includes assigning a column graph which represents an ordering relationship between columns of the first table, the column graph being a row graph from a second table. Finally, the process includes interpreting the table, row graph or column graph as a finite sheaf (see specification beginning at page 15, line 14 to page 16, line 2 for example).

Although not limited to the specific example of Figures 3 and 4A, these figures are nevertheless used as an illustration for understanding the terminology used in Claim 43, for example. The triangle of Figure 3 may be described in its constituent components in a table (44 in Figure 4A). The row graph is the set of numbered circles with interconnecting arrows

(46) shown in Figure 4A, though ordering of the different rows are externally specified (not generated according to the method). As can be seen, the different rows in table 44 are arranged in the row graph 46 according to the externally specified ordering relationship between the rows of the table. Another step according to the method of Claim 43 is defining a finite distributive lattice (FDL) from the row graph to have a member for each row and each distinct combination of the rows of the first table. In contrast, Altschuler describes defining a table that contains both entity rows and relationship rows, and therefore does not create a one-to-one correspondence between rows and entities of the particular type. Furthermore, in contrast to Claim 43, Altschuler provides an ordering relationship that is computed from the table as part of the method, while Claim 43 specifically requires the ordering relationship be externally specified.

Altschuler is asserted in the outstanding Office Action as anticipating Claim 43. Applicant respectfully traverses the rejection.

Altschuler is directed to a system for using lattice structures in which data is represented by graphs in which data are modeled as graphs having "entity" nodes which depict properties or attributes of the data and "relationship" links which denote relationship among the entities (column 13, lines 33-36). Various types of data may be represented using annotated entity relationship diagrams (or "a-ERDs"). (column 13, lines 39-44). Figure 47 is a table data structure 4700 which represents usage log graph 800 of Figure 8. The data used in Altschuler are represented as graphs in an extended sense, in that relationship links can attach to other relationship links (as shown in Figure 8).

Comparing amended Claim 43 to Altschuler, amended Claim 43 is directed to entering attribute data into rows of a first table, and specifying a row graph which represents an externally specified ordering relationship between the rows of the first table. In contrast, Altschuler describes defining a table that contains both entity rows and relationship rows, and

therefore does not create a one-to-one correspondence between rows and entities of the particular type. Furthermore, in contrast to Claim 43, Altschuler provides an ordering relationship that is computed from the table as part of the method, while Claim 43 specifically requires the ordering relationship be externally specified. Claim 43 further requires defining a finite distributive lattice from the row graph to have a member for each row and each combination of the rows of the first table. In contrast, Altschuler is not directed to a row-based system. Rather, Altschuler is directed to a system that operates on annotated entity relationship diagrams (a-ERD) that comprise an entire table of data (e.g., Figure 47 as an example of an a-ERD). Altschuler does not disclose defining a finite distributive lattice from a row graph to have a member for each row and for each distinct combination of rows of the first table. Thus, while amended Claim 43 is directed to specific features within a table, Altschuler is directed to something entirely different, namely providing relationships between different annotated entity relationship diagrams. Furthermore, Claim 43 requires that the row graph represent an externally specified ordering relationship between the rows of the first table. In contrast, a feature of Altschuler is that the relationship between different a-ERDs is determined automatically and not externally specified.

Claim 43 also requires the assignment of a column graph which represents an ordering relationship between columns of the first table, the column graph being a row graph of a second table. Altschuler neither teaches nor suggests this feature. Likewise, Claim 43 requires interpreting the table, row graph or column graph as a finite sheaf. In contrast Altschuler provides no column graph and the columns are fixed so as to always have the same columns. In contrast to Altschuler, the 1-2-1 correspondence between the columns of table and attributes of the entity type, make it such that each table will have its own set of columns according to the present invention. Once again, Altschuler neither teaches nor suggest this feature.

Consequently, it is respectfully submitted that Claim 43, as amended, is patentably distinct over Altschuler. Although Claims 45, 50, 56, 67-72, 74 and 80, are of differing scope and/or statutory subject class, it is respectfully submitted that each of these claims also patentably defines over Altschuler for at least the same reasons discussed above with regard to Claim 43.

Turning to the rejections of pending claims over the combination of Altschuler in view of Oles, each of these rejections is based on the assertion of Altschuler as applied to the independent claims discussed above. Oles is asserted for its alleged description of defining a set of algebraic operators to operate on graphs, and lattice data models. However, Oles does not disclose the features of Claim 43 as discussed above, namely the step of defining a finite sheaf from a column graph in particular to have a member for each row and each distinct combination of the rows of the first table, as an example. As a consequence, no matter how Oles is combined with Altschuler, the combination would neither teach nor suggest the features of independent Claim 43, and therefore the combination does not teach or suggest the features of dependent Claims 49, 51-55, 57-66, 73, 75-79, and 81-90.

Consequently, in view of the present amendment and in light of the foregoing comments, it is respectfully submitted that the invention defined by Claims 1-43, 45-67 and 69-90, as amended, is patentably distinguishing over the prior art. The present application is therefore believed to be in condition for formal allowance and an early and favorable reconsideration of this application is therefore requested.

Respectfully submitted,

OBLON, SPIVAK, McCLELLAND,
MAIER & NEUSTADT, P.C.

Bradley D. Lytle
Attorney of Record
Registration No. 40,073

Customer Number
22850

Tel: (703) 413-3000
Fax: (703) 413 -2220
(OSMMN 06/04)

OBLON, SPIVAK, ET AL
DOCKET #: 208485US-8
INV: DAVID M. BUTLER
SERIAL NO: 09/852,781
Response to Office Action Dated May 20, 2004
Annotated Sheet

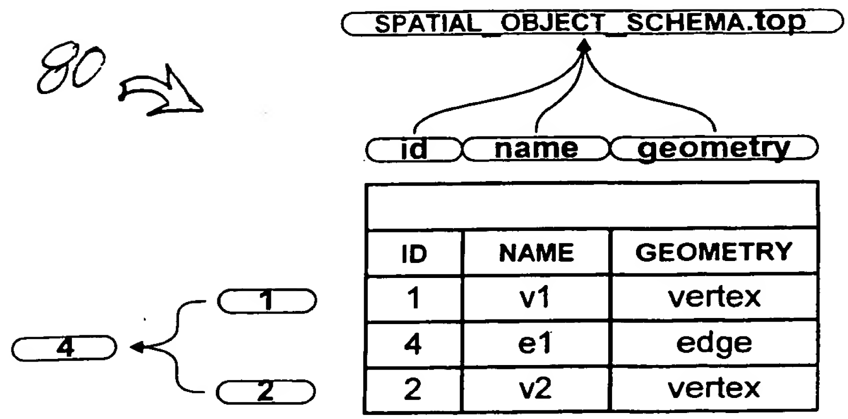

ID	NAME	GEOMETRY
1	v1	vertex
4	e1	edge
2	v2	vertex

Figure 4b

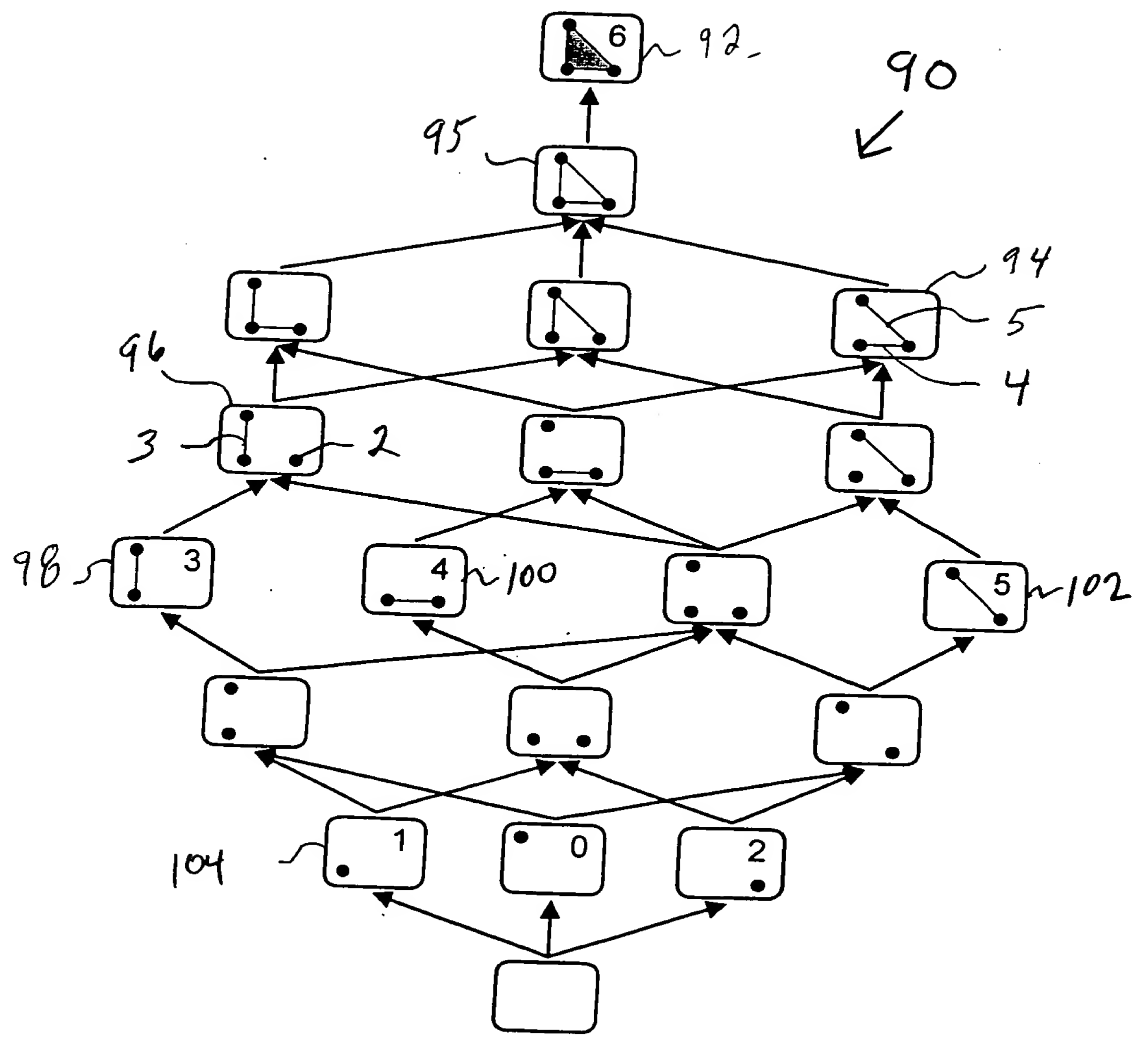

Figure 4c